

Subspace Clustering Based Tag Sharing for Inductive Tag Matrix Refinement with Complex Errors

Yuqing Hou[†] Zhouchen Lin^{*†} Jin-ge Yao[§]

[†]Key Lab. of Machine Perception (MOE), School of EECS, Peking University, P. R. China

[§]Institute of Computer Science and Technology, Peking University, P. R. China
{yqh, zlin, yaojinge}@pku.edu.cn

ABSTRACT

Annotating images with tags is useful for indexing and retrieving images. However, many available annotation data include missing or inaccurate annotations. In this paper, we propose an image annotation framework which sequentially performs tag completion and refinement. We utilize the subspace property of data via sparse subspace clustering for tag completion. Then we propose a novel matrix completion model for tag refinement, integrating visual correlation, semantic correlation and the novel studied property of complex errors. The proposed method outperforms the state-of-the-art approaches on multiple benchmark datasets even when they contain certain levels of annotation noise.

Keywords

image annotation; subspace clustering; matrix completion; complex errors

1. INTRODUCTION

It is useful to annotate images with textual tags for the purpose of image indexing and retrieval. To annotate proper tags, one need to bridge the gap between low level visual features of an image and corresponding high level semantic information [16]. Since manual annotation is labor intensive, automatic annotation has aroused much attention. Many machine learning based approaches have been developed.

Currently many image annotation data have been collected from crowdsourcing services [21, 12], providing large amount of data for training while being noisy due to annotation errors. Annotation errors are usually complex and mainly come in two forms: missing tags and inaccurate tags. Most image annotation approaches solely focus on one of those two, either trying to impute the missing tags (tag completion/tag assignment) [11] or correcting inaccurate tags (tag refinement) [22, 8, 16]. Other existing methods fail to model the complex errors properly. They either treat them in the same way [24], ignoring the complex property of the

errors, or rigidly assign fixed weights to different kinds of errors [11], having no adaptability when working on different datasets with different levels of annotation errors.

In this paper, we propose a framework called Subspace clustering and Matrix completion with Complex errors (SMC). Since current tag refinement methods suffer from the extreme sparsity problem [20], SMC performs tag completion and refinement sequentially. During tag completion, SMC tries to introduce many additional proper tags to images via exploring subspace property in the image collection. We then adapt the inductive matrix completion [13] model to perform the following tag refinement procedure, utilizing side information such as the correlation between visual features and their corresponding tags (visual correlation), correlation between the semantic information of tags (semantic correlation) and the complex errors.

The main contributions of this paper include:

- We perform tag completion and tag refinement sequentially, showing that tag refinement benefits from tag completion.
- We formulate tag completion in a subspace clustering framework to tackle the extreme sparsity problem.
- We novelly adapt the inductive matrix completion model for tag refinement, taking visual correlation, semantic correlation and our novelly studied complex errors property into consideration.

2. THE PROPOSED FRAMEWORK

2.1 Overview

We denote the observed tag matrix as $\mathbf{O} \in \{0, 1\}^{N_i \times N_t}$, where each row corresponds to one image, each column corresponds to one textual tag, and N_i and N_t denote the number of images and tags, respectively. \mathbf{O}_{ij} takes value 1 only if image i is annotated with tag j and 0 otherwise.

We are targeting at modifying the values in matrix \mathbf{O} by matrix completion methods to perform image annotation. After the matrix completion procedure, if the value of \mathbf{O}_{ij} changes from nonzero (zero) to zero (nonzero), we say that the algorithm removes (adds) tag j from (to) image i . Methods based on matrix completion are robust and efficient since they only operate on the tag matrix, avoiding error propagation from image segmentation.

However, in many cases \mathbf{O} is so sparse that some columns have at most one known entries and some rows have no known entries at all, making existing methods not applicable [20]. In order to overcome such extreme sparsity, we first

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '16, July 17-21, 2016, Pisa, Italy

© 2016 ACM. ISBN 978-1-4503-4069-4/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2911451.2914693>

perform tag completion to make \mathbf{O} denser, creating a better condition for the following tag refinement procedure. More specifically, we perform subspace clustering over images and share tags within subspaces.

For tag refinement, existing methods usually depends heavily on image segmentation and visual feature extraction accuracies [1]. However, image segmentation and feature extraction procedures always contain a lot of noises, which affect the following annotation procedure severely. Meanwhile, recent matrix completion based methods [10, 8, 9] stand out due to their robustness and efficiency, since these algorithms avoid the image segmentation procedure.

Our proposed framework is called Subspace clustering and Matrix completion with Complex errors (SMC), because it utilizes the subspace property of image collections (Section 2.2) and addresses the complex errors and side information in an inductive matrix completion model for tag refinement (Section 2.3).

2.2 Subspace Clustering for Tag Completion

2.2.1 Subspace Clustering

It is reasonable to assume that images belonging to different categories are approximately sampled from a mixture of several low-dimensional subspaces. The membership of the data points to the subspaces is unobserved, leading to the challenging problem of subspace clustering. Here the goal is to cluster data into k clusters with each cluster corresponding to a subspace.

One of the state-of-the-art method is the sparse subspace clustering (SSC) model [5]. The idea behind SSC is to express a data point as a linear (or affine) combination of neighboring data points. The neighbors can be any other points in the data set. While every point is a combination of all other data points, SSC seeks for the sparsest representation among all the candidates by minimizing the number of nonzero coefficients [6].

We denote the set of images, represented as visual feature vectors, as $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n]$. Assuming that they are drawn from a union of k subspaces. Each column of \mathbf{V} can be represented by a linear combination of the bases in a "dictionary". SSC uses the matrix \mathbf{V} itself as the dictionary while explicitly considering noise:

$$\min_{\mathbf{Z}, \mathbf{E}} \quad \|\mathbf{Z}\|_1 + \mu \|\mathbf{E}\|_F^2, \quad (1)$$

$$\text{s.t.} \quad \mathbf{V} = \mathbf{VZ}^\top + \mathbf{E}, \text{diag}(\mathbf{Z}) = \mathbf{0}, \mathbf{Z}\mathbf{1} = \mathbf{1}, \quad (2)$$

where $\mathbf{Z}^\top = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n]$ is the coefficient matrix with each \mathbf{z}_i being the representation of \mathbf{v}_i and \mathbf{E} is the error matrix. This problem can be solved efficiently using modern sparse optimization algorithms, such as linearized alternating direction methods [17].

Given a sparse representation for each data point, we can define the affinity matrix as $\mathbf{A} = |\mathbf{Z}| + |\mathbf{Z}^\top|$. Subspaces are then obtained by applying spectral clustering to the Laplacian matrix of \mathbf{A} [5].

2.2.2 Tag Sharing

We improve the search based neighbor voting algorithm proposed in [18] to share tags in each cluster separately. We rank all the tags for the cluster, taking tag frequency, tag co-occurrence and local frequency into consideration. The elements of tag matrix after tag sharing are no longer binary

but take values in $[0, 1]$, representing the confidence level between each image-tag pair.

2.3 Matrix Completion for Tag Refinement

The tag completion procedure makes the tag matrix much denser and thus avoids the extreme sparsity problem. Then we can refine the tag matrix. In our framework we novelly adapt the inductive matrix completion model (IMC) [13] for tag refinement, due to its scalability and capability of incorporating various kinds of side information.

2.3.1 Inductive Matrix Completion

Let $\mathbf{v}_i \in \mathbb{R}^{f_i}$ denote the f_i -dimensional feature vector of image i and $\mathbf{t}_j \in \mathbb{R}^{f_t}$ denote the f_t -dimensional feature vector of tag j . Let $\mathbf{V} \in \mathbb{R}^{N_i \times f_i}$ denote the feature matrix of N_i images, where the i -th row is the image feature vector \mathbf{v}_i^\top , and $\mathbf{T} \in \mathbb{R}^{N_t \times f_t}$ denote the feature matrix of N_t tags, where the i -th row is the tag feature \mathbf{t}_i^\top .

For image annotation, we assume that the tag matrix can be approximated by applying visual feature vectors and tag feature vectors associated with its row and column entries onto an underlying low-rank matrix \mathbf{M} , i.e. $\mathbf{O} \approx \mathbf{VMT}^\top$, where $\mathbf{M} = \mathbf{PQ}^\top$ [13] and $\mathbf{P} \in \mathbb{R}^{f_i \times r}$ and $\mathbf{Q} \in \mathbb{R}^{r \times f_t}$ are of rank $r \ll N_i, N_t$. The goal is to solve the following problem:

$$\min_{\mathbf{P}, \mathbf{Q}} \quad \text{loss}(\mathbf{O}, \mathbf{VPQ}^\top \mathbf{T}^\top) + \lambda_1 (\text{rank}(\mathbf{PQ}^\top)). \quad (3)$$

A common choice for the loss function is the squared loss. The low-rank constraint on \mathbf{PQ}^\top makes (3) NP-hard. A standard relaxation is to use the trace norm, i.e. sum of singular values. Minimizing the trace-norm of $\mathbf{M} = \mathbf{PQ}^\top$ is equivalent to minimizing $\frac{1}{2}(\|\mathbf{P}\|_F^2 + \|\mathbf{Q}\|_F^2)$ [13]. The relaxed optimization problem we use in this work is therefore:

$$\min_{\mathbf{P}, \mathbf{Q}} \quad \|\mathbf{O} - \mathbf{VPQ}^\top \mathbf{T}^\top\|_F^2 + \frac{\lambda_1}{2} (\|\mathbf{P}\|_F^2 + \|\mathbf{Q}\|_F^2). \quad (4)$$

2.3.2 Visual Correlation

We want to get the refined tag matrix $\hat{\mathbf{O}} = \mathbf{VPQ}^\top \mathbf{T}^\top$ from the original tag matrix \mathbf{O} . Here we represent the i th row of $\hat{\mathbf{O}}$ as $\hat{\mathbf{O}}_i$, corresponding to the refined tag vector of image i . Thus we can measure the correlation between image i and image j in two ways: 1) similarity between image features \mathbf{v}_i and \mathbf{v}_j , 2) similarity between refined tag vectors $\hat{\mathbf{O}}_i$ and $\hat{\mathbf{O}}_j$. Since visually similar images often belong to similar themes and thus are annotated with similar tags, these two kinds of similarities should be correlated.

Such visual correlation can be enforced by solving the following optimization

$$\min_{\mathbf{P}, \mathbf{Q}} \quad \sum_{i=1}^{N_t} \sum_{j=1}^{N_t} \|\hat{\mathbf{O}}_i - \hat{\mathbf{O}}_j\|^2 g_{ij}, \quad (5)$$

where $\|\hat{\mathbf{O}}_i - \hat{\mathbf{O}}_j\|^2$ measures the similarity between tag vectors $\hat{\mathbf{O}}_i$ and $\hat{\mathbf{O}}_j$ and g_{ij} measures the similarity between visual features \mathbf{v}_i and \mathbf{v}_j . In this work, we adopt cosine similarity, i.e. $g_{ij} = \cos(\mathbf{v}_i, \mathbf{v}_j)$. The formulation forces tag vectors with large similarities also have large similarity in their corresponding visual features and vice versa.

The formulation can be rewritten as

$$\min_{\mathbf{P}, \mathbf{Q}} \quad \text{Tr}(\hat{\mathbf{O}} \mathbf{L}_v \hat{\mathbf{O}}^\top) = \min_{\mathbf{P}, \mathbf{Q}} \quad \text{Tr}(\mathbf{VPQ}^\top \mathbf{T}^\top \mathbf{L}_v \mathbf{TQ}^\top \mathbf{V}^\top), \quad (6)$$

where $\mathbf{L}_v = \text{diag}(\mathbf{G}\mathbf{1}) - \mathbf{G}$ is the Graph Laplacian [3] of the similarity matrix $\mathbf{G} = (g_{ij})$.

2.3.3 Semantic Correlation

Similarly, we can also enforce semantic correlation between tags. Since each column of the matrix $\hat{\mathbf{O}}$ represents the feature of a tag, we can measure the correlation between two tags using the similarity between their corresponding column vectors of $\hat{\mathbf{O}}$. Meanwhile, semantic similarity between two tags can be measured using word vectors. These two kinds of similarities should be correlated as well.

We can enforce the semantic correlation by solving the following optimization, in a similar form as (6):

$$\min_{\mathbf{P}, \mathbf{Q}} \text{Tr}(\hat{\mathbf{O}}^\top \mathbf{L}_s \hat{\mathbf{O}}) = \min_{\mathbf{P}, \mathbf{Q}} \text{Tr}(\mathbf{TQP}^\top \mathbf{V}^\top \mathbf{L}_s \mathbf{VPQ}^\top \mathbf{T}^\top), \quad (7)$$

where $\mathbf{L}_s = \text{diag}(\mathbf{H}\mathbf{1}) - \mathbf{H}$ is the Graph Laplacian of the similarity matrix $\mathbf{H} = (h_{ij})$, with each element $h_{ij} = \cos(\mathbf{t}_i, \mathbf{t}_j)$.

2.3.4 Features Vectors

We utilize DeCAF₆ [4] to extract 4,096-dimensional visual features for each image, which have high level information. Meanwhile, we adopt pre-trained word embedding vectors (word2vec) [19] to construct 300-dimensional features for each tag, trying to capture semantic information.

2.3.5 Complex Errors

As we have mentioned, annotation errors come in two forms: missing tags and inaccurate tags. Since human beings are relatively reasonable, the user-provided tags are reasonably accurate to certain level [24]. Users might miss one or several proper tags among the few related tags, but may become less probable to add one or several inaccurate tags from the massive unrelated tag sets [12]. In other words, if an image is not originally annotated with a tag, it is more likely that they really have no relation at all. Thus the errors are mainly composed of inaccurate tags rather than missing tags. And we should pay more attention to denoise the inaccurate tags rather than completing the missing ones.

To model the complex structure of errors, we improve the matrix completion model by putting less weights on the unannotated positions:

$$\min_{\mathbf{P}, \mathbf{Q}} \|\mathbf{O} - \mathbf{VPQ}^\top \mathbf{T}^\top\|_F^2 - \mu \|\mathbf{U}_\Omega(\mathbf{O} - \mathbf{VPQ}^\top \mathbf{T}^\top)\|_F^2, \quad (8)$$

where Ω represents the positions where the images are originally not annotated. \mathbf{U} is a projection operator and μ acts as a weighting parameter which changes adaptively in different datasets according to their noise levels.

Existing methods never model these two kinds of errors separately. They simply model the errors as Laplacian noise [24] or Gaussian noise [22]. To our knowledge, our model is the first to model the missing errors and inaccurate errors separately. The model can further adapt to different datasets according to their noise levels.

3. FINAL MODEL

Based on the components regarding low-rankness, visual correlation, semantic correlation and complex errors, we formulate the objective function as follows:

$$\begin{aligned} & \min_{\mathbf{P}, \mathbf{Q}} \|\mathbf{O} - \mathbf{VPQ}^\top \mathbf{T}^\top\|_F^2 - \mu \|\mathbf{U}_\Omega(\mathbf{O} - \mathbf{VPQ}^\top \mathbf{T}^\top)\|_F^2 \\ & + \frac{\lambda_1}{2} (\|\mathbf{P}\|_F^2 + \|\mathbf{Q}\|_F^2) + \\ & \lambda_2 [\text{Tr}(\mathbf{VPQ}^\top \mathbf{T}^\top \mathbf{L}_v \mathbf{TQP}^\top \mathbf{V}^\top + \mathbf{TQP}^\top \mathbf{V}^\top \mathbf{L}_s \mathbf{VPQ}^\top \mathbf{T}^\top)]. \end{aligned}$$

By solving \mathbf{P} and \mathbf{Q} we can then construct the refined tag matrix $\hat{\mathbf{O}} = \mathbf{VPQ}^\top \mathbf{T}^\top$ and use it for refined annotation.

We set the regularization terms of visual correlation and semantic correlation with the same weight λ_2 for simplicity. This simplification does not harm performance, as we find during preliminary experiments.

This objective function is non-convex. To solve the optimization problem, we adapt the solver for low-rank empirical risk minimization for multi-label learning (LEML) [23], which naturally fits for the settings of large-scale multi-label learning with missing labels. The solver uses alternating minimization (fix \mathbf{P} and solve for \mathbf{Q} and vice versa) to update the variables. When either \mathbf{P} or \mathbf{Q} is fixed, the resulting subproblem in one variable (\mathbf{Q} or \mathbf{P}) can be solved using iterative conjugate gradient procedure.

4. EXPERIMENTAL EVALUATION

4.1 Datasets and Experimental Setup

We evaluate our proposed SMC framework on two benchmark datasets: Labelme [21] and MIRFlickr-25K [12]. Table 1 demonstrates the detailed statistics. These two datasets, especially MIRFlickr-25K, are rather noisy, with a number of the tags being misspelled or meaningless. Hence, a pre-processing procedure is performed. We match each tag with entries in the Wikipedia thesaurus and only retain the tags in accordance with Wikipedia.

Table 1: Statistics of 2 Datasets

Statistics	Labelme	MIRFlickr-25K
No. of images	2,900	25,000
Vocabulary Size	495	1,386
Tags per Image (mean/max)	10.5/48	12.7/76
Images per Tag (mean/max)	67.1/379	416.5/76,890

We compare our method with the state-of-the-art methods, including matrix completion-based models (i.e. LRES [24], TCMR [8], RKML [9]), search-based models (i.e. JEC [18], TagProp [11], and TagRelevance [15]), mixture models (i.e. CMRM [14] and MBRM [7]) and co-regularized learning model (FastTag [2]). The tag refinement procedure by itself, denoted as SMC_IMC, is also compared to verify the benefit from the tag completion procedure. We tune the parameters on the validation set of the two datasets separately for every method in comparison. Note that the weighting parameter μ we tune changes from 0.4 (Labelme) to 0.7 (MIRFlickr-25K), confirming that as the data become more and more noisy, we should pay more attention to the noisy tags and less on missing tags.

We measure all the methods in terms of *average precision@N* ($AP@N$) and *average recall@N* ($AR@N$). In the top N completed tags, *precision@N* is to measure the ratio of correct tags and *recall@N* is to measure the ratio of missing ground-truth tags, both averaged over all test images.

4.2 Evaluation and Observation

Table 2 and Table 3 show performance comparisons on the two datasets, respectively.

We can observe that: 1) Generally, methods achieve better performance on Labelme, since tags in MIRFlickr-25K are more noisy. 2) Methods based on matrix completion, such as SMC, LRES and TCMR, usually achieve the best performances. 3) Our SMC framework shows increasing advantage to LRES as the data become more and more noisy, justifying our assumption and model on the noises. 4) SMC nearly

Table 2: Performance Comparison on Labelme

	Labelme					
	N = 2		N = 5		N = 10	
	AP	AR	AP	AR	AP	AR
SMC	0.51	0.36	0.46	0.50	0.35	0.62
SMC_IMC	0.47	0.34	0.40	0.48	0.31	0.59
LRES [24]	0.42	0.32	0.35	0.45	0.27	0.56
TCMR [8]	0.44	0.32	0.37	0.45	0.29	0.55
RKML [9]	0.21	0.14	0.19	0.20	0.14	0.22
JEC [18]	0.33	0.29	0.27	0.38	0.20	0.48
TagProp [11]	0.39	0.31	0.33	0.45	0.25	0.56
TagRel [15]	0.43	0.32	0.34	0.45	0.27	0.55
CMRM [14]	0.20	0.14	0.18	0.19	0.12	0.22
MBRM [7]	0.23	0.14	0.18	0.20	0.12	0.27
FastTag [2]	0.43	0.34	0.37	0.44	0.28	0.57

Table 3: Performance Comparison on MIRFlickr-25K

	MIRFlickr-25K					
	N = 2		N = 5		N = 10	
	AP	AR	AP	AR	AP	AR
SMC	0.53	0.39	0.40	0.47	0.33	0.61
SMC_IMC	0.45	0.34	0.36	0.43	0.30	0.52
LRES [24]	0.43	0.35	0.32	0.40	0.26	0.45
TCMR [8]	0.45	0.35	0.35	0.41	0.28	0.48
RKML [9]	0.21	0.15	0.13	0.23	0.13	0.22
JEC [18]	0.33	0.30	0.25	0.34	0.19	0.35
TagProp [11]	0.39	0.35	0.28	0.37	0.20	0.41
TagRel [15]	0.42	0.34	0.30	0.37	0.20	0.40
CMRM [14]	0.20	0.15	0.13	0.18	0.11	0.20
MBRM [7]	0.22	0.16	0.13	0.18	0.10	0.22
FastTag [2]	0.43	0.35	0.30	0.41	0.27	0.42

outperforms all the other algorithms in all cases. 5) Performance comparison between SMC and SMC_IMC demonstrate the remarkable benefit of tag completion for tag refinement. 6) Performance on MIRFlickr-25K in some sense provides an evidence for the robustness of SMC.

5. CONCLUSION

In this work we present an effective framework for image annotation by performing tag completion and tag refinement sequentially. Our method first clusters images using sparse subspace clustering and shares tags using a neighbor voting algorithm, then refines tags by adapting inductive matrix completion while novelly utilizing visual and semantic information. Experiments show the effectiveness of our framework and suggest that tag refinement can benefit a lot from performing tag completion first.

Acknowledgments

Zhouchen Lin is supported by National Basic Research Program of China (973 Program) (grant no. 2015CB352502), National Natural Science Foundation (NSF) of China (grant nos. 61272341 and 61231002), and Microsoft Research Asia Collaborative Research Program.

6. REFERENCES

- [1] G. Carneiro, A. Chan, P. Moreno, and N. Vasconcelos. Supervised learning of semantic classes for image annotation and retrieval. *TPAMI*, 2007.
- [2] M. Chen, A. Zheng, and K. Weinberger. Fast image tagging. In *ICML*, 2013.
- [3] F. Chung. *Spectral graph theory*. 1997.
- [4] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. *arXiv preprint arXiv:1310.1531*, 2013.
- [5] E. Elhamifar and R. Vidal. Sparse subspace clustering. In *CVPR*, 2009.
- [6] E. Elhamifar and R. Vidal. Clustering disjoint subspaces via sparse representation. In *ICASSP*, 2010.
- [7] S. Feng, R. Manmatha, and V. Lavrenko. Multiple Bernoulli relevance models for image and video annotation. In *CVPR*, 2004.
- [8] Z. Feng, S. Feng, R. Jin, and A. Jain. Image tag completion by noisy matrix recovery. In *ECCV*. 2014.
- [9] Z. Feng, R. Jin, and A. Jain. Large-scale image annotation by efficient and robust kernel metric learning. In *ICCV*, 2013.
- [10] A. Goldberg, B. Recht, J. Xu, R. Nowak, and X. Zhu. Transduction with matrix completion: Three birds with one stone. In *NIPS*, 2010.
- [11] M. Guillaumin, T. Mensink, J. Verbeek, and C. Schmid. Tagprop: Discriminative metric learning in nearest neighbor models for image auto-annotation. In *ICCV*, 2009.
- [12] M. Huiskes and M. Lew. The MIR Flickr retrieval evaluation. In *ICMIR*, 2008.
- [13] P. Jain and I. Dhillon. Provable inductive matrix completion. *arXiv preprint arXiv:1306.0626*, 2013.
- [14] J. Jeon, V. Lavrenko, and R. Manmatha. Automatic image annotation and retrieval using cross-media relevance models. In *SIGIR*, 2003.
- [15] X. Li, C. Snoek, and M. Worring. Learning social tag relevance by neighbor voting. *TM*, 2009.
- [16] X. Li, T. Uricchio, L. Ballan, M. Bertini, C. G. Snoek, and A. Del Bimbo. Socializing the semantic gap: A comparative survey on image tag assignment, refinement and retrieval. *arXiv preprint arXiv:1503.08248*, 2015.
- [17] Z. Lin, R. Liu, and Z. Su. Linearized alternating direction method with adaptive penalty for low-rank representation. In *NIPS*, 2011.
- [18] A. Makadia, V. Pavlovic, and S. Kumar. A new baseline for image annotation. In *ECCV*. 2008.
- [19] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [20] N. Natarajan and I. Dhillon. Inductive matrix completion for predicting gene-disease associations. *Bioinformatics*, 2014.
- [21] B. Russell, A. Torralba, K. Murphy, and W. Freeman. Labelme: a database and web-based tool for image annotation. *IJCV*, 2008.
- [22] L. Wu, R. Jin, and A. Jain. Tag completion for image retrieval. *TPAMI*, 2013.
- [23] H. Yu, P. Jain, P. Kar, and I. Dhillon. Large-scale multi-label learning with missing labels. In *ICML*, 2014.
- [24] G. Zhu, S. Yan, and Y. Ma. Image tag refinement towards low-rank, content-tag prior and error sparsity. In *ACM MM*, 2010.